

A Noise-Shaping Coder Topology for 15 + Bit Converters

H03M3/04

L. RICHARD CARLEY, MEMBER, IEEE

E P. 267-273

Abstract—A novel topology for high-precision noise-shaping converters that can be integrated on a standard digital IC process is presented. This topology employs a multibit noise-shaping coder and a novel form of dynamic element matching to achieve high accuracy and long-term stability without requiring precision matching of components. A fourth-order noise-shaping D/A conversion system employing a 3-bit quantizer and a dynamic element matching internal D/A converter, fabricated in a standard double-metal 3- μ m CMOS process, achieved 16-bit dynamic range and a harmonic distortion below -90 dB. This multibit noise-shaping D/A conversion system achieved performance comparable to that of a 1-bit noise-shaping D/A conversion system that operated at nearly four times its clock rate.

I. INTRODUCTION

ONE-BIT noise-shaping converters, referred to as either delta-sigma modulators ($\Delta\Sigma$'s) or sigma-delta modulators, have recently achieved popularity for use in integrated circuit data converters, both A/D [1] and D/A [2]. Their attractiveness for IC systems that incorporate digital filtering and signal processing with A/D and/or D/A conversion is, in part, due to the fact that they employ a 1-bit internal D/A converter which does not require precision component matching. They can be implemented using a standard digital CMOS process without the economically costly addition of precision thin-film resistors or the use of laser trimming. However, experience has revealed that high integral linearity is very difficult to achieve, because it is not possible to avoid overloading the 1-bit loop quantizer. One obvious solution is to use a multibit quantizer in the loop employing sufficient levels to prevent quantizer overload. Eliminating quantizer overload, in addition to improving integral linearity, also facilitates the design of higher-order (> 2) noise-shaping loops, because the low-frequency oscillations observed in higher-order $\Delta\Sigma$ loops [1] are a result of quantizer overload [7]. For these reasons, a multibit noise-shaping conversion system can achieve a signal-to-noise ratio (SNR) comparable to that of a second-order 1-bit noise-shaping ($\Delta\Sigma$)

D/A conversion system operating at a much higher sampling rate. For example, the prototype noise-shaping D/A conversion system presented in this paper, which operated at 3.2 MHz and employed a 3-bit quantizer, achieved performance comparable to that of a $\Delta\Sigma$ -type conversion system that operates at 11.3 MHz [2]. The lower clock rate is particularly important since dynamic power consumption in the CMOS logic increases as the square of the clock frequency.

As stated above, many advantages result from using a multibit noise-shaping converter instead of a $\Delta\Sigma$ converter. However, there is one major disadvantage: the integral linearity of the noise-shaping conversion system is no better than the integral linearity of the multibit internal D/A converter [3], [4]. Achieving high integral linearity, hence low total harmonic distortion (THD), normally requires precision component matching. In this case, multibit noise-shaping converters cannot be fabricated using an inexpensive digital CMOS process, hence their relative lack of use as part of digital signal processing IC systems to date.

This paper presents a multibit noise-shaping coder topology that incorporates an internal D/A converter employing a novel form of "dynamic element matching" [5], [6] to achieve excellent integral and differential linearity, while requiring only modest component matching. For example, the prototype IC D/A conversion system has a peak element mismatch of nearly 0.2 percent, and a measured peak integral linearity error of only 0.0022 percent of full scale. This new topology allows multibit noise-shaping coders to employ internal D/A converters which achieve high linearity without requiring component trimming. Therefore, by employing a dynamic element matching internal D/A converter, multibit noise-shaping coders can be integrated with digital signal processing electronics using an inexpensive digital CMOS process as demonstrated by the prototype system.

II. NOISE-SHAPING CODER TOPOLOGY

Fig. 1 shows the topology for a multibit noise-shaping A/D conversion system and Fig. 2 shows the D/A conversion system. In the case of an A/D conversion system,

Manuscript received August 2, 1988; revised December 6, 1988. This work was supported in part by the National Science Foundation under Grant ENG-8451496 and by a grant from Analog Devices Corporation. The author is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213.
IEEE Log Number 8826153.




Fig. 1. Topology of noise-shaping A/D conversion systems. OV is the oversampling ratio and f_s is the sampling rate at the system's digital output. The number of bits at the output L is a function of the decimating filter architecture and the SNR of the noise-shaping coder.

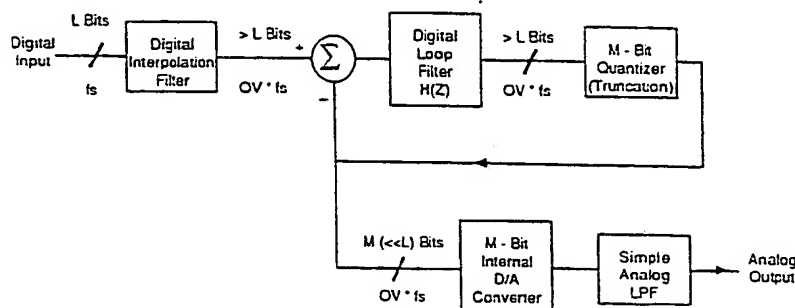


Fig. 2. Topology of noise-shaping D/A conversion systems. OV is the oversampling ratio and f_s is the sampling rate at the system's digital input. The number of bits at the interpolating filter output is greater than L due to arithmetic operations in the filter, and is therefore a function of the interpolating filter architecture.

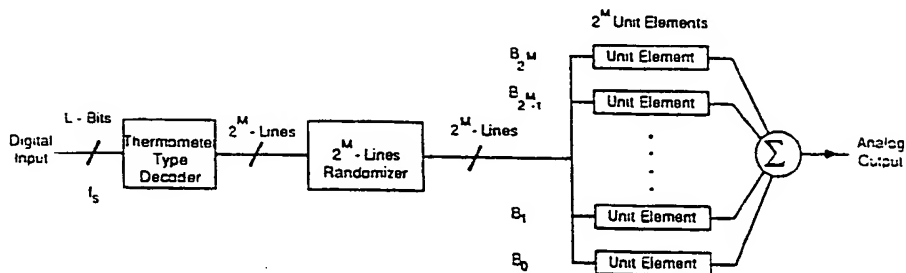


Fig. 3. Topology of dynamic element matching internal D/A converter. A thermometer-type decoder sets the number of output lines high that is equal to the digital input.

the quantizer is a true A/D converter—normally a bank of comparators, while in the case of a D/A conversion system, the quantizer merely corresponds to truncation. In both cases, a multibit internal D/A converter is necessary.

Assuming that quantizer overload does not occur, the design of multibit noise-shaping loops is quite simple compared to the design of $\Delta\Sigma$ loops, because the gain of the quantizer is known (1 LSB/digital level), the quantization error is bounded between $\pm 1/2$ and $-1/2$ LSB, and the quantization error in this case can be modeled as an additive noise that is independent of the input signal [12], [13]. Nonlinear numerical techniques have been applied to optimize the conversion system's performance for a given internal D/A converter through the choice of the loop-filter pole and zero locations [13].

Because the quantizer error is bounded, a worst-case bound can be computed for the range signal at the quantizer's input. The worst-case noise gain is determined by assuming that the quantization error will take on a se-

quence of worst-case values, $+1/2$ or $-1/2$ LSB [12], [13]. The noise gain can be computed by summing the magnitude of the impulse response of the transfer function from the quantizer output back to its input. The quantizer must have enough levels to code the worst-case noise at its input plus the input signal range without overload. For example, if the $H(z)$ selected gives a worst-case noise gain of 4, then the noise due to quantization at the quantizer's input will be ± 2 LSB's. If the input signal also occupied ± 2 LSB's of range then the quantizer and the internal D/A converter would need eight levels.

III. INTERNAL D/A CONVERTER TOPOLOGY

The dynamic element matching D/A converter topology consists of two parts: a parallel unit-element D/A converter structure and a digital "randomizer" which controls the connections between the input and the parallel D/A converter. First we will consider the parallel D/A con-

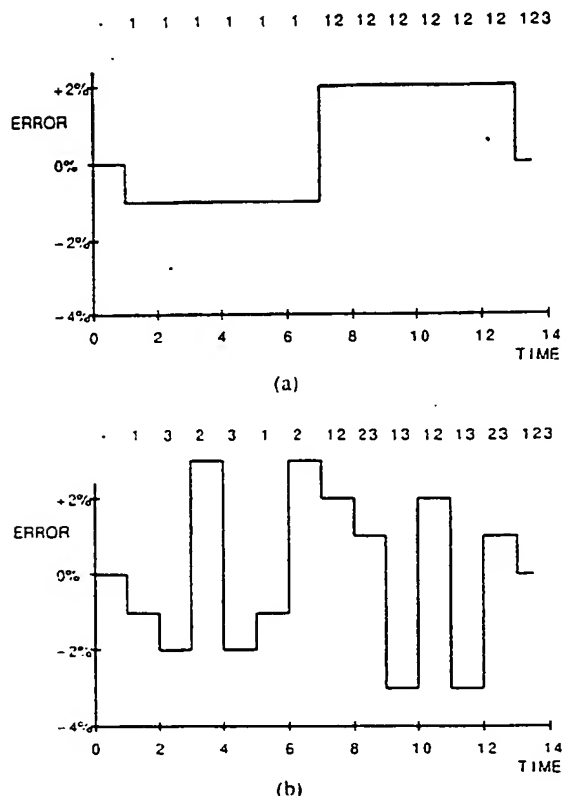


Fig. 4. How dynamic element matching works. (a) The error of a 2-bit parallel D/A converter implemented as the sum of three elements. (b) The error of the same D/A converter with the randomizer implementing dynamic element matching. Element 1 is 1 percent below desired value, element 2 is 3 percent above desired value, and element 3 is 2 percent below the desired value. The numbers above each trace indicate which elements are active. The input is 0, 1, 1, 1, 1, 1, 2, 2, 2, 2, 2, 2, 3.

verter and the performance achievable by employing dynamic element matching in this case. Then we will consider possible implementations of the randomizer.

A. Dynamic Element Matching Approach

The dynamic element matching D/A converter topology can be constructed using any D/A converter in which the N th output level is generated by activating N approximately equal-valued elements, typically resistors or capacitors, and summing up their charge or current or voltage (see Fig. 3). Dynamic element matching is implemented by choosing different elements to represent the N th level as a function of time. The "randomizer" block varies which elements will be used to represent the N th level on each clock cycle (see Fig. 4). The goal of this approach is to convert the error due to element mismatch from a dc offset into an ac signal of equivalent power which, in an oversampling converter, can be partially removed by filtering. In Fig. 4, even when the input is constant, the error is a wide-band noise signal. With ideal randomization, a mismatch between the unit elements would be converted into a white-noise signal with zero mean error and a variance equal to the root-mean-square (rms) error between the

individual unit elements. In an oversampling data conversion system, which typically operates at oversampling ratios of 64:1 or more, nearly all of the error power can be filtered out.

First, let us consider the linearity of this D/A converter. For a dc input code of N , each element is active, on average, N out of every M clock cycles, where M is the total number of elements. Therefore, each element of the D/A converter acts individually as a duty-cycle modulator and the integral linearity is limited only by the product of the fractional element mismatch ($\Delta E/E$) and the fractional clock jitter ($\Delta T/T$) [5], [6]. A second practical limit on the integral linearity results because there is normally a small change in the charge (or current) transferred by each element as a function of the number of elements active. With careful choice of D/A converter topology and the use of a precision clock, extremely high dc integral linearity can be achieved, even when the elements match very poorly. However, the element mismatch now appears as an ac noise signal added to the D/A converter output. The constraint on the element matching has been converted from a constraint based on the converter's linearity to a constraint based on the converter dynamic range.

If small scale factor errors are ignored, the maximum noise signal $n(t)$ varies in a parabolic fashion from zero at either zero or full scale to a maximum at half of full scale. At this maximum, $n(t)$, relative to the internal D/A converter's full scale M , is

$$\text{rms} \left[\frac{n(t)}{M} \right] = \frac{\text{rms} \left[\frac{\Delta E}{E} \right]}{2\sqrt{M}}$$

Note, this expression is a factor of $\sqrt{2}$ smaller than the sum of $M/2$ random element errors because of the assumption that scale factor errors are negligible.

In an oversampling converter, only noise power in the signal passband is important. Assuming an ideal randomizer, the noise signal will be white and the rms fraction of the internal D/A converter's full scale for the noise signal in the passband ($n_p(t)$) will be

$$\text{rms} \left[\frac{n_p(t)}{M} \right] = \frac{\text{rms} \left[\frac{\Delta E}{E} \right]}{2\sqrt{M \cdot OF}}$$

where OF is the oversampling ratio. Note, although the noise-shaping A/D conversion system shapes quantization noise in order to achieve a greater than 3-dB increase in SNR per octave, the internal D/A element mismatch noise is not in the feedback loop (see Fig. 2) and cannot be affected by feedback.

Since an R -bit D/A converter would require $2^R - 1$ elements, only a limited number of bits can be generated using this topology. The unusual nature of noise-shaping coders, that resolution is gained by oversampling the signal




Fig. 5. Topology of the butterfly-type randomizer. Example of a three-stage eight-line butterfly randomizer. Each pair of switches marked with the same label is controlled to either exchange the two signal lines or pass them directly to the next stage.

IV. DESIGN EXAMPLE

In order to validate the predictions for the linearity and SNR of the new noise-shaping coder topology, an experimental fourth-order noise-shaping D/A conversion system has been designed, optimized using CLANS [13], simulated, fabricated, and tested. Experimental versions using both resistor unit elements and capacitor unit elements in the internal D/A converter have been studied.

A. System Topology

Fig. 2 shows the block diagram of the experimental noise-shaping D/A conversion system. Interpolation, the process of increasing the sampling rate, was performed in two stages. First, the input signal was interpolated by a factor of 2 and filtered using a standard 64-tap low-pass FIR filter topology. Then, the FIR filter output was interpolated by a factor of 32 and filtered using a comb filter [8]. The combined filter attenuates frequencies above 30 kHz by more than 80 dB with a 3.2-MHz clock. Extensive simulations indicate that the SNR at the output of the noise-shaping loop with this interpolator architecture is less than 3 dB below the SNR with an ideal interpolating filter. Although the filter complexity could be reduced still further, a dramatic decrease in the loop's SNR was observed for filter attenuations below 80 dB.

The noise-shaping loop filter ($H(z)$) in Fig. 2 employs a fourth-order filter with four poles at $z = 1$ and three zeros positioned to maintain system stability and to minimize the worst-case noise gain of the loop [9], [10]. CLANS [13] was used to optimize the locations of the three zeros in order to maximize the loop's SNR. Maximum tolerable noise gain is a constraint used in choosing the location of the zeros of the loop filter. A maximum noise gain of 4 was chosen as the constraint when placing the closed-loop poles for the noise-shaping loop of the prototype system. Therefore, the quantization error at the noise-shaping loop's output, which is in the range of $+1/2$ to $-1/2$ LSB, is amplified in the worst case by a factor of 4 at the quantizer input. Hence, the noise signal at the quantizer input will be in the range of $+2$ to -2 LSB. In order to prevent the noise-shaping quantizer from being overloaded, the 18-bit input signal is scaled so that its full-scale range is between $+2$ and -2 LSB, resulting in a total range of eight levels or 3 bits.

B. 3-bit Internal D/A Converter

The capacitor version of the dynamic element matching 3-bit internal D/A converter is shown in Fig. 6. The D/A converter's N th output level is generated by charging all capacitors to a $+5$ -V reference level and then switching N of them into the summing junction of the I^- operational amplifier and the rest into the summing junction of the I^+ amplifier during each clock cycle. Note that by always switching all of the capacitors between the same two voltages on every clock cycle many potential complications

processing but that linearity must be provided by the D/A converter, makes them a perfect application for this type of dynamic element matching which is limited to only a few bits but can provide extremely high accuracy. In addition, only a few levels, typically two to four, are necessary to prevent quantizer overload in noise-shaping coders; therefore, the number of parallel elements is not prohibitive for this application.

B. The Randomizer

The randomizer connects the M outputs from the decoder to the M switching elements in a time-varying fashion. The number of possible connections is M factorial. Therefore, when M is small (on the order of 3 or 4) it is possible to randomly select between all possible connections. However, when M is large (e.g., 8) the number of possible connections is so large that it may be necessary to select a subset of connections in order to conserve die area. For example, an ideal eight-level randomizer, one which connects each of the eight inputs to eight outputs, would have to include 40 320 possible connections.

One simple approach to randomizing over a subset of possible connections would be to have an M -port barrel shifter which rotates one increment after each clock. This represents only M of the M factorial possible permutations. This approach would completely decorrelate successive output errors only if the mismatch between elements were independent of the element's position on the die. Unfortunately, just the opposite is typically true. Adjacent elements are normally much more likely to match than distant elements due to gradients in process parameters across the wafer.

A compromise between these two extremes is the "butterfly" randomizer proposed by Kenney [14]. The butterfly randomizer circuit consists of a series of butterfly networks coupling the inputs to the outputs (see Fig. 5). In order that any input can be connected to any output, the number of butterfly stages should be at least equal to the number of bits in the internal D/A converter. More butterfly stages can be added if it is necessary to cover a larger fraction of possible connections. A pseudorandom sequence generator would normally be used to generate the random control sequences for the butterfly switches [15]–[17].

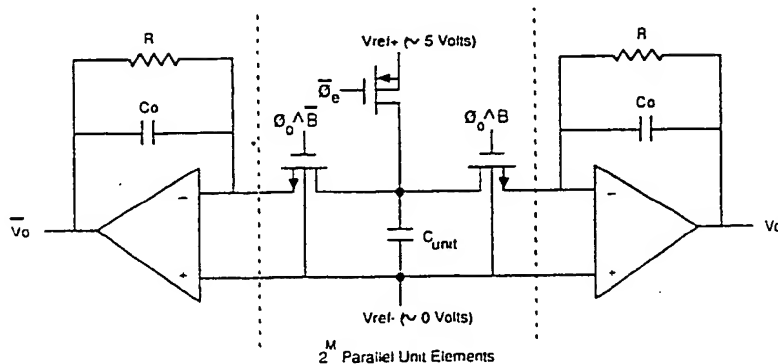


Fig. 6. Topology of the capacitive 3-bit internal D/A converter. ϕ_e and ϕ_o are the even and odd switch clocks, respectively, and B is the digital input to that element. The unit C was 2 pF.

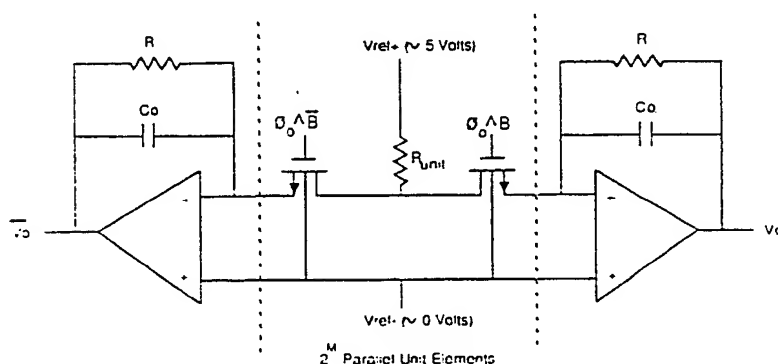


Fig. 7. Topology of the resistive 3-bit internal D/A converter. ϕ_e and ϕ_o are the even and odd switch clocks, respectively, and B is the digital input to that element. The unit R was 14K.




Fig. 8. Photomicrograph of randomizer and 3-bit capacitive D/A converter.

are avoided: varying current draw on the reference source, complications due to dielectric absorption in the capacitors, and even failure of the capacitor voltage to fully settle to either the reference voltage or ground. The primary nonlinearity error of this topology comes from changes in the transient at the op-amp summing junction with the number of elements selected.

A resistor version of the 3-bit D/A converter was also fabricated (see Fig. 7). As with the capacitor version, the N th output level was generated by switching N resistors into the summing junction of the I^- operational amplifier and the rest into the summing junction of the I^+ operational amplifier. The other end of each resistor was tied to the reference voltage. Note that even and odd switch

BEST AVAILABLE COPY




Fig. 9. Plot of SNR as a function of input amplitude. In this case distortion is included with the noise. Input signal is a 1-kHz sine wave.

clocks were used, even in the resistive case, to prevent varying propagation delays in the computation of B from changing the ON period of the element.

Dynamic element matching was implemented by randomizing the elements activated by each code over time. Eight unit elements were used rather than the minimum of seven. At least one element was always OFF. The choice between using 2^M elements and $2^M - 1$ unit elements in the internal D/A converter is arbitrary, and is not important to its operation. In this case, eight levels were selected to make the design of the butterfly randomizer symmetric. The randomizer circuit consisted of a series of three "butterfly" networks coupling the inputs to the output (see Fig. 6). This randomizer implements 4096 of the available 40320 combinations, about one-tenth of the possible connections. Simulations verified that this partial randomization resulted in good decorrelation of internal D/A converter errors in the face of both linear and quadratic variations in element value with position. Fig. 8 shows a photomicrograph of the capacitive 3-bit D/A converter, the decoder, and the digital randomizer circuitry.

The analog output from the internal 3-bit D/A converter (see Fig. 2) was filtered by a fifth-order Butterworth filter. The real axis pole was provided by the first op amp and the two complex pole pairs were implemented as a cascade of two Sallen-Key second-order sections. Because the SNR and the open-loop gain of the operational amplifiers limit the measurements of the D/A converter's performance, the analog filter was implemented off-chip using low-noise operational amplifiers.

Although all of the prototype system was not fabricated in IC form, other researchers have fabricated digital interpolating filters, digital noise-shaping loops, and analog output filters in CMOS processes equivalent in performance to those required in this system [2], [11]. For example, a higher-order (1:256) interpolating filter and a second-order noise-shaping loop were implemented by Naus *et al.* in approximately 18 mm² of die area [2]. Therefore, we conclude that by testing the novel portion of this




Fig. 10. PSD at D/A converter output. The input is a full-scale 1-kHz signal.




Fig. 11. PSD at analog filter output with a 0-dB 1-kHz input signal. The input is a full-scale 1-kHz signal. Note, the amplitude of the input signal has been artificially reduced by passing it through a 1-kHz notch filter in order to decrease the dynamic range to allow accurate measurement of the PSD.

system, the dynamic element matching internal D/A converter, in IC form we have proven that integration of the entire system is viable.

C. Results

In order to be compatible with a double-metal digital CMOS process, the capacitive version of the internal D/A converter used the field oxide between metal-1 and metal-2 as the unit capacitor. Unfortunately, this oxide was extremely irregular, the capacitor mismatch was over 1 percent. Although the integral linearity was excellent as expected, the noise resulting from element mismatch was unacceptably large. Therefore, the measured results reported in the rest of this section are for the resistive version of the internal D/A converter. The resistors in the resistive version of the internal D/A converter were 3- $\mu\text{m} \times 2.7\text{-mm}$ polysilicon lines with a resistance of about 14K. On selected die rms element variations as low as 0.05 percent were observed.

Fig. 9 shows the SNR at the analog filter output as a function of the input amplitude for a 1-kHz sinusoidal input. The system achieves a dynamic range of about 98

dB. Fig. 10 shows the power spectral density (PSD) of the 3-bit D/A converter output without the analog filter for a 0-dB 1-kHz sinusoidal input. The flattening of the PSD of the noise at approximately 1 MHz is a result of the zeros incorporated into the noise-shaping loop's transfer function to control the noise gain. Note that this signal contains noise components from two sources: quantization noise from the noise-shaping loop and the approximately "white" noise resulting from the randomized element mismatch.

Fig. 11 shows the PSD at the output of the analog filter. Simulations with an ideal 3-bit D/A converter indicate that the PSD of the in-band noise is dominated, in the audio frequency range, by the component mismatch noise rather than the quantization noise from the noise-shaping loop. The dynamic range limited by the quantization alone was 108 dB for this noise-shaping coder. The harmonic distortion components are each more than 96 dB below the maximum input level.

V. CONCLUSIONS

A new topology for multibit high-precision noise-shaping converters that incorporates a dynamic element matching internal D/A converter has been presented. This new topology makes possible the construction of multibit noise-shaping coders in a standard digital CMOS process. Experimental verification has been provided via a prototype 16-bit oversampling D/A conversion system that uses an all-digital noise-shaping loop followed by a dynamic element matching 3-bit D/A converter that was fabricated in 3- μ m CMOS. The 16-bit D/A conversion system achieved a peak dynamic range of 98 dB and a THD of approximately -94 dB at a sampling rate of only 3.2 MHz. This represents almost a factor of 4 decrease in clock rate from a $\Delta\Sigma$ -type conversion system of comparable performance [2] and a potential decrease of nearly 16-fold in the dynamic power consumption of the CMOS logic implementing the interpolating filter and noise-shaping loop.

The multibit nature of the noise-shaping loop, which both increases the initial accuracy of the quantizer and facilitates the use of higher-order noise-shaping coders, makes it possible to attain performance equivalent to that of $\Delta\Sigma$ converters operating at much higher clock rates. And it is the dynamic element matching technique which makes possible multibit noise-shaping coders that can be incorporated into a standard digital CMOS process with digital filters and signal processors.

ACKNOWLEDGMENT

The author would like to thank J. Kenney and C. Wolff for their help in laying out and testing several versions of the experimental D/A converter reported herein.

REFERENCES

- [1] J. C. Candy, "A use of double integration in sigma-delta modulation," *IEEE Trans. Commun.*, vol. COM-33, pp. 249-258, Mar. 1985.
- [2] P. J. A. Naus et al., "A CMOS stereo 16-bit D/A converter for digital audio," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 3, pp. 390-395, June 1987.
- [3] J. W. Scott, W. Lee, C. Giancarlo, and C. G. Sodini, "A CMOS slope adaptive delta modulator," in *ISSCC Dig. Tech. Papers*, Feb. 1986, pp. 130-131.
- [4] R. W. Adams, "Design and implementation of an audio 18-bit A/D converter using oversampling techniques," *J. Audio Eng. Soc.*, vol. 34, no. 3, pp. 153-166, Mar. 1986.
- [5] R. J. Van De Plassche, "A monolithic 14-bit D/A converter," *IEEE J. Solid-State Circuits*, vol. SC-14, no. 3, pp. 552-556, June 1979.
- [6] K. B. Klaassen, "Digitally controlled absolute voltage division," *IEEE Trans. Instrum. Meas.*, vol. IM-24, no. 3, pp. 106-112, June 1975.
- [7] C. Wolff and L. R. Carley, "Modeling the quantizer in higher-order delta-sigma modulators," presented at the Int. Symp. Circuits Syst., Helsinki Finland, June 1988.
- [8] S. Chu and C. S. Burrus, "Multirate filter designs using comb filters," *IEEE Trans. Circuits Syst.*, vol. CAS-31, pp. 913-924, Nov. 1984.
- [9] J. Kenney, "Design methodology for N -th order noise-shaping coders," Masters Project Rep., Dept. of ECE, Carnegie-Mellon Univ., Pittsburgh PA, Jan. 1988.
- [10] B. P. Agrawal and K. Shenoi, "Design methodology for $\Sigma\Delta$ M," *IEEE Trans. Commun.*, vol. COM-31, pp. 360-369, Mar. 1983.
- [11] Y. Matsuya et al., "A 16-bit oversampling a-to-d conversion technology using triple-integration noise shaping," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 6, pp. 921-929, Dec. 1987.
- [12] L. R. Carley, "An oversampling analog-to-digital converter topology for high resolution signal acquisition systems," *IEEE Trans. Circuits Syst.*, vol. CAS-34, no. 1, pp. 83-91, Jan. 1987.
- [13] J. Kenney and L. R. Carley, "CLANS: A high-level synthesis tool for high resolution data converters," in *Proc. 1988 IEEE Int. Conf. Computer-Aided Des.* (Santa Clara, CA), Nov. 1988, pp. 496-499.
- [14] L. R. Carley and J. Kenney, "A 16-bit 4th order noise-shaping D/A converter," in *Proc. 1988 Custom Integrated Circuits Conf.* (Rochester, NY), May 1988.
- [15] J. H. Lindholm, "An analysis of the pseudo-randomness properties of subsequences of long m -sequences," *IEEE Trans. Inform. Theory*, vol. IT-14, pp. 569-567, July 1968.
- [16] J. L. Manos, "Some techniques for testing pseudo-random number sequences," Lincoln Labs., Lexington, MA, Tech. Note 1974-44, 1974.
- [17] A. C. Davies, "Properties of waveforms obtained by nonrecursive digital filtering of pseudo-random binary sequences," *IEEE Trans. Computers*, vol. C-20, pp. 270-281, Mar. 1971.



L. Richard Carley (S'77-M'84) received the S.B., S.M., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1976, 1978, and 1984, respectively.

In 1984 he joined Carnegie Mellon University, Pittsburgh, PA, as an Assistant Professor of Electrical and Computer Engineering. His research is in the area of analysis, automated synthesis, and simulation of mixed analog/digital systems.

Dr. Carley received a National Science Foundation Presidential Young Investigator Award in 1985, and a Best Paper Award at the 1987 Design Automation Conference in Miami.

THIS PAGE BLANK (USPTO)